

CLAIMS

What is claimed is:

1. A method comprising:

trapping, by a processor, a change in execution among schedulable entities; and
tracking an execution of a schedulable entity that is being switched in for
execution as a result of the change in execution.

2. The method of claim 1, wherein the tracking is performed by a privileged entity
and further comprising:

calculating, by the privileged entity, an estimated resource requirement for the
schedulable entity that is being switched in for execution from the tracking of a previous
execution of the schedulable entity; and

scheduling, by the privileged entity, the schedulable entity that is being switched
in for execution according to its estimated resource requirement.

3. The method of claim 2, wherein the privileged entity is a virtual machine monitor
and the schedulable entities are selected from the group consisting of processes, threads,
operating systems, virtual machines, and child virtual machine monitors.

4. The method of claim 2, wherein the virtual machine monitor comprises:

an idle detector to receive notice from the processor of the change in execution
and to derive a measured value for a schedulable entity that is being switched out of
execution;

a proportional integral derivative (PID) controller logically coupled to the idle
detector to receive the measured value and to calculate the estimated resource
requirement required by the schedulable entity that is being switched out of execution;

8 and

9 a scheduler logically coupled to the PID controller to receive the estimated
10 resource requirement and to determine a schedule of execution for the schedulable entity
11 that is being switched out of execution.

1 5. The method of claim 2, wherein calculating an estimated resource requirement
2 comprises:

3 assigning an initial value as the estimated resource requirement for the
4 schedulable entity that is being switched in for execution;

5 reducing the estimated resource requirement for the schedulable entity if the
6 schedulable entity completes execution before the estimated resource requirement is
7 exhausted; and

8 increasing the estimated resource requirement for the schedulable entity if the
9 schedulable entity does not complete execution before the estimated resource requirement
10 is exhausted.

1 6. The method of claim 2 further comprising:

2 initiating, by the privileged entity, the change in execution.

1 7. The method of claim 6, wherein the tracking of the execution is performed as part
2 of the change in execution initiated by the privileged entity.

1 8. The method of claim 2 further comprising:

2 initiating, by the processor, the change in execution if the change in execution is
3 being requested by the privileged entity.

1 9. The method of claim 1, wherein trapping a change in execution comprises:

2 detecting an instruction to change a state register that identifies a schedulable

3 entity.

1 10. The method of claim 9 further comprising:
2 comparing, by the processor, the state register that identifies the schedulable
3 entity being switched in for execution with a state match register that identifies a
4 schedulable entity that is to be tracked, wherein the schedulable entity being switched
5 into execution is tracked by the processor if the state register and the state match register
6 match.

1 11. The method of claim 1, wherein trapping a change in execution comprises:
2 detecting an instruction to change between privileged and non-privileged modes.

1 12. The method of claim 1, wherein the schedulable entities are selected from the
2 group consisting of operating system processes, operating system threads, virtual
3 machines, and instruction streams to be executed by the processor.

1 13. A machine-readable medium providing instructions, which when executed by a
2 machine, causes the machine to perform operations comprising:
3 trapping, by a processor, a change in execution among schedulable entities; and
4 tracking an execution of a schedulable entity that is being switched in for
5 execution as a result of the change in execution.

1 14. The machine-readable medium of claim 13, wherein the tracking is performed by
2 a privileged entity and further comprising:
3 calculating, by the privileged entity, an estimated resource requirement for the
4 schedulable entity that is being switched in for execution based on the tracking of a
5 previous execution of the schedulable entity; and

6 scheduling, by the privileged entity, the schedulable entity that is being switched
7 in for execution according to the estimated resource requirement.

1 15. The machine-readable medium of claim 14, wherein the privileged entity is a
2 virtual machine monitor and the schedulable entities are selected from the group
3 consisting of processes, threads, operating systems, virtual machines, and child virtual
4 machine monitors.

1 16. The machine-readable medium of claim 14, wherein the virtual machine monitor
2 comprises:

3 an idle detector to receive notice from the processor of the change in execution
4 and to derive a measured value for a schedulable entity that is being switched out of
5 execution;

6 a proportional integral derivative (PID) controller logically coupled to the idle
7 detector to receive the measured value and to calculate the estimated resource
8 requirement required by the schedulable entity that is being switched out of execution;
9 and

10 a scheduler logically coupled to the PID controller to receive the estimated
11 resource requirement and to determine a schedule of execution for the schedulable entity
12 that is being switched out of execution.

1 17. The machine-readable medium of claim 14, wherein calculating an estimated
2 resource requirement comprises:

3 assigning an initial value as the estimated resource requirement for the
4 schedulable entity that is being switched in for execution;

5 reducing the estimated resource requirement for the schedulable entity if the
6 schedulable entity completes execution before the estimated resource requirement is
7 exhausted; and

8 increasing the estimated resource requirement for the schedulable entity if the
9 schedulable entity does not complete execution before the estimated resource requirement
10 is exhausted.

1 18. The machine-readable medium of claim 14 further comprising:
2 initiating, by the privileged entity, the change in execution.

1 19. The machine-readable medium of claim 18, wherein the tracking of the execution
2 is performed as part of the change in execution initiated by the privileged entity.

1 20. The machine-readable medium of claim 14 further comprising:
2 initiating, by the processor, the change in execution if the change in execution is
3 being requested by the privileged entity.

1 21. The machine-readable medium of claim 13, wherein trapping a change in
2 execution comprises:
3 detecting an instruction to change a state register that identifies a schedulable
4 entity.

1 22. The machine-readable medium of claim 13 further comprising:
2 comparing, by the processor, the state register that identifies the schedulable
3 entity being switched in for execution with a state match register that identifies a
4 schedulable entity that is to be tracked, wherein the schedulable entity being switched
5 into execution is tracked by the processor if the state register and the state match register
6 match.

1 23. The machine-readable medium of claim 13, wherein trapping a change in

2 execution comprises:
3 detecting an instruction to change between privileged and non-privileged modes.

1 24. The machine-readable medium of claim 13, wherein the schedulable entities are
2 selected from the group consisting of operating system processes, operating system
3 threads, virtual machines, and instruction streams to be executed by the processor.

1 25. An apparatus comprising:
2 a memory;
3 a processing unit coupled to the memory and configured to trap to a privileged
4 entity, a change in execution among schedulable entities; and
5 the privileged entity executed from the memory to cause the processing unit to
6 track an execution of a schedulable entity that is being switched in for execution as a
7 result of the change in execution.

1 26. The apparatus of claim 25, wherein the privileged entity further causes the
2 processing unit to calculate an estimated resource requirement for the schedulable entity
3 that is being switched in for execution based on the tracking of a previous execution of
4 the schedulable entity and to schedule the schedulable entity that is being switched in for
5 execution according to the estimated resource requirement.

1 27. The apparatus of claim 26, wherein the privileged entity is a virtual machine
2 monitor and the schedulable entities are selected from the group consisting of processes,
3 threads, operating systems, virtual machines, and child virtual machine monitors.

1 28. The apparatus of claim 27, wherein the virtual machine monitor comprises:
2 an idle detector to receive notice from the processing unit of the change in

3 execution and to derive a measured value for a schedulable entity that is being switched
4 out of execution;

5 a proportional integral derivative (PID) controller logically coupled to the idle
6 detector to receive the measured value and to calculate the estimated resource
7 requirement for the schedulable entity that is being switched out of execution; and

8 a scheduler logically coupled to the PID controller to receive the estimated
9 resource requirement and to determine a schedule of execution for the schedulable entity
10 that is being switched out of execution.

1 29. The apparatus of claim 26, wherein the privileged entity further causes the
2 processing unit to calculate an estimated machine resource requirement by:

3 assigning an initial value as the estimated resource requirement for the
4 schedulable entity;

5 reducing the estimated resource requirement for the schedulable entity if the
6 schedulable entity completes execution before the estimated resource requirement is
7 exhausted; and

8 increasing the estimated resource requirement for the schedulable entity if the
9 schedulable entity does not complete execution before the estimated resource requirement
10 is exhausted.

1 30. The apparatus of claim 25, wherein the processing unit is further configured to
2 trap a change in execution by detecting an instruction to change a state register that
3 identifies a schedulable entity.

1 31. The apparatus of claim 25, wherein the processing unit is further configured to
2 trap a change in execution by detecting an instruction to change between privilege and
3 non-privilege modes.

